Chapter 4

Implementing Land Surface Data Assimilation

In this Chapter we describe in detail how the one-dimensional land-surface model of Chapter 3 is used in a fully four-dimensional (space and time) data assimilation algorithm. We also match the general estimation formulation of Chapter 2 with the variables of the hydrologic model.

A key assumption in the hydrologic model is to neglect lateral flow in the unsaturated zone, which is reasonable for terrain with moderate relief and on the spatial scales under consideration. The model domain thus breaks down into a collection of one-dimensional vertical cells or pixels (Figure 4.1). The horizontal resolution of the land surface model is determined by the availability of the micro-meteorologic data and of the land surface parameters, such as land cover and soil texture maps, and by computational resources. Since we aim to estimate the land surface states at this resolution, we refer to this scale as the *estimation pixels*.

From a simulation perspective, the estimation pixels are uncoupled but for larger-scale patterns of the inputs such as the meteorologic forcings as well as the land cover and soil texture classes. Lateral unsaturated moisture and heat fluxes between different pixels are neglected. When the model is incorporated into the data assimilation algorithm, however, the pixels are also coupled through the statistics of the uncertain inputs. The model error, representing for example errors in surface forcings, and the initial condition parameters in different pixels are assumed to be random fields which are correlated over time and space. The assimilation algorithm is therefore fully four-dimensional (space and time).

The horizontal resolution of the brightness images is solely determined by the sensor. We call this scale the *observation pixels*. In the foreseeable future, the brightness images will only be available at a resolution much lower than the resolution of the land surface model, that is each observation pixel generally contains several estimation pixels. We therefore need to develop the assimilation algorithm together with a general downscaling capability.

This Chapter is organized as follows. We first describe the model domain and the spatial discretization (Section 4.1). Next, we define the state vector for the land surface application (Section 4.2). The initial condition parameterization is explained in Section 4.3, the state equation is defined in Section 4.4, and the model error terms are identified in Section 4.5. Section 4.7 contains the definition of the assimilated data and the measurement operator. We also present the general downscaling scheme, which is based on an appropriate definition of the measurement operator. Next, Section 4.8 reviews the covariance models

of the uncertain inputs that we use in the assimilation. After a brief discussion of how the adjoint and the tangent-linear models are derived (Section 4.9), we conclude this Chapter with a description of the temporal discretization of the model equations (Section 4.10).

4.1 Model Domain and Spatial Discretization

The model domain for the estimation of soil moisture and temperature consists of $N_{ep} \equiv N_x \cdot N_y$ vertical columns or *estimation pixels* (Figure 4.1). Typical scales for the size of the estimation pixels are $1 \ldots 10km$. Each column has N_z vertical nodes for the saturation. The vertical resolution is typically six layers ($N_z = 7$). Starting at the surface, the thickness increases downwards ($5cm$, $10cm$, $15cm$, $15cm$, $15cm$, and $30cm$). For the soil surface temperature there is only one node per column, likewise for the canopy temperature, the interception storage, the canopy air temperature, and the canopy vapor pressure.

In the unsaturated zone, horizontal moisture and heat fluxes over the scale of an estimation pixel are insignificant compared to vertical fluxes. We therefore neglect horizontal unsaturated fluxes, which makes the model very computationally efficient. Even though the columns are uncoupled as far as lateral moisture and energy exchange is concerned, they are in fact connected through horizontal correlations of the micro-meteorologic inputs and the soil and vegetation parameters. Meteorologic and geologic processes generally vary over a scale larger than the scale of an estimation pixel for soil moisture. Such correlations provide the large scale patterns observed by remote sensing instruments. Within the data assimilation algorithm, the estimation pixels are also coupled through the covariance structure of the uncertain inputs (Section 4.8).

4.2 State Vector

The state variables for all pixels are collected into the state vectors X and Y (4.1). The canopy temperature T_c, the vapor pressure e_a and the temperature T_a in the canopy air space, which are subject to algebraic equations, are concatenated into the state vector X. Likewise, the saturation W_g, the soil temperature T_g, and the interception water content W_c, which obey ordinary differential equations, are collected in the vector Y. The superscript in parentheses indicates the number of the estimation pixel.

$$X = \begin{bmatrix} T_c^{(1)}/T_0 \\ e_a^{(1)}/e_{20} \\ T_a^{(1)}/T_0 \\ \cdots\cdots \\ T_c^{(2)}/T_0 \\ e_a^{(2)}/e_{20} \\ T_a^{(2)}/T_0 \\ \cdots\cdots \\ \vdots \end{bmatrix} \qquad Y = \begin{bmatrix} \begin{pmatrix} W_{g1}^{(1)} \\ \vdots \\ W_{gN_z}^{(1)} \end{pmatrix} \\ T_g^{(1)}/T_0 \\ W_c^{(1)}/(W_c^{\max,(1)} + \gamma_s) \\ \cdots\cdots\cdots\cdots\cdots\cdots \\ \vdots \end{bmatrix} \tag{4.1}$$

All temperatures are scaled with the reference temperature $T_0 = 273.15K$. The vapor pressure e_a in the canopy air space is scaled with the water vapor pressure $e_{20} = 23.4mb$ at $20°C$. The interception storage is dimensionless, but needs to be scaled to be of order

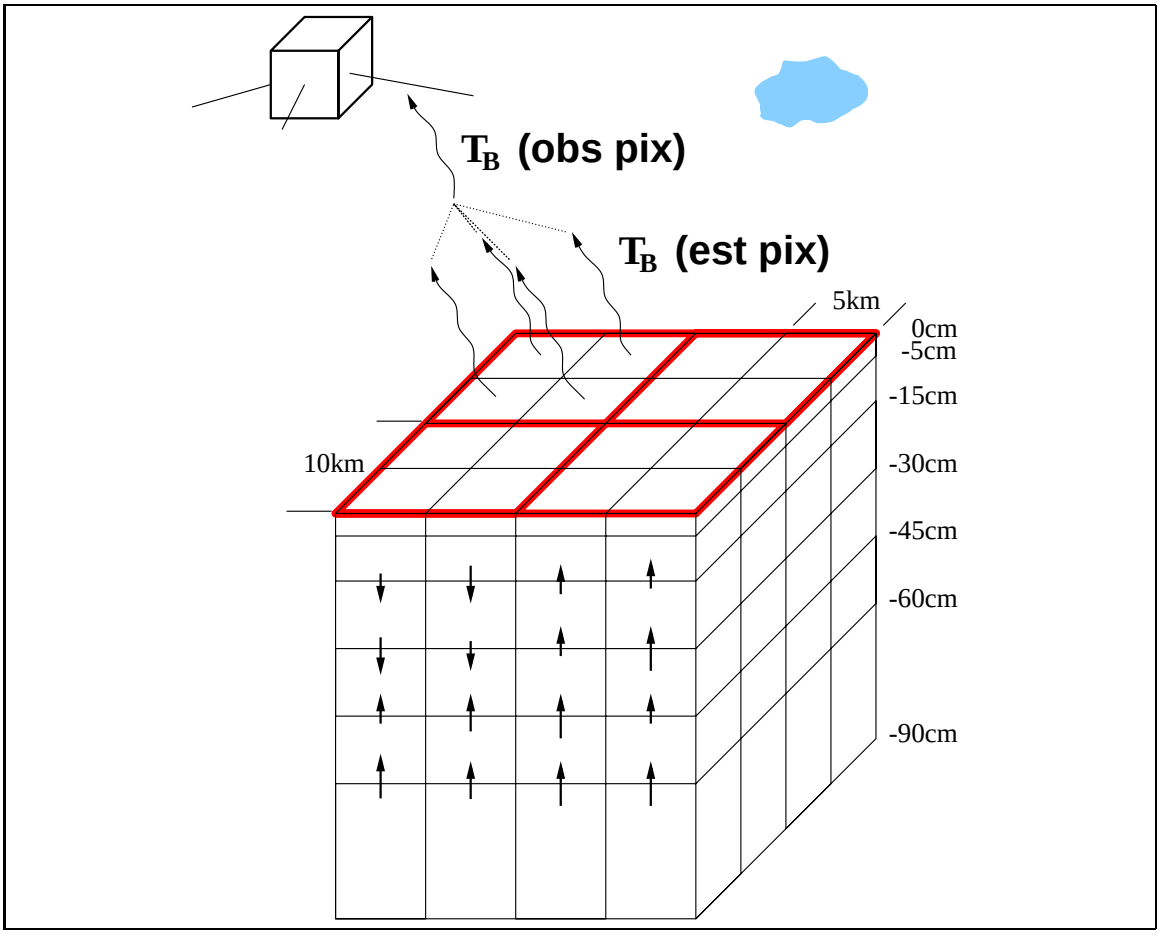

Figure 4.1: Schematic of the model grid with typical length scales. Note the huge ratio of horizontal to vertical scales. In the unsaturated zone, vertical fluxes dominate. Lateral unsaturated fluxes of moisture and energy are neglected (Section 4.1). Soil moisture estimates are obtained on the $5km$-scale of the estimation pixels (fine outline). Remotely sensed brightness measurements are available on the $10km$-scale of the observation pixels (thick outline). The downscaling is achieved by including in the measurement operator the computation of the average of the brightness temperature over an observation pixel. For L-band brightness temperatures, the arithmetic average is appropriate [Drusch et al., 1999]. This average is then subtracted from the data to provide the data misfit. By respecting the information contained in the fine-scale model inputs, the algorithm implicitly distributes the innovation from the measurement over all estimation pixels within a given observation pixel (Section 4.7).

one. This is done with $(W_c^{\max} + \gamma_s)$, where $W_c^{\max}$ depends on the pixel (3.16) and γ_s is a constant scaling factor. The saturation W_g is already dimensionless and on the order of one. It does not require further scaling.

The time-dependent vector $Y(t)$ is of length $N_Y = N_x \cdot N_y \cdot (N_z + 2)$. The length of $X(t)$ is $N_X = N_x \cdot N_y \cdot 3$. Altogether, there are $N_x \cdot N_y \cdot (N_z + 5)$ states at every time step. For a two-week period with $15min$ time steps, we need $N_t \approx 1300$ time steps. Using $N_x = 16$, $N_y = 32$, and $N_z = 7$, the total number of states is on the order of 10^7.

4.3 Initial Condition

In this Section, we match the land-surface variables to the initial condition parameterization $(2.1a)$.

4.3.1 Initial Saturation

The parameterization for the initial saturation $W_g(t = 0)$ maps the entire profile of each pixel onto a single uncertain parameter γ_g, the transformed total initial water storage, as follows. First, we decompose the initial saturation profile of each pixel into the total initial water storage in the column W_g^{store} and the shape of the initial profile W_g^{shape}, which is a vector of length N_z.

$$W_{gi}(t = 0) = W_g^{\text{store}} \cdot W_{gi}^{\text{shape}} \tag{4.2}$$

The total initial water storage is defined as

$$W_g^{\text{store}} = \sum_{i=1}^{N_z} \bar{\Delta}_i \theta_{si} W_{gi}(t = 0) \tag{4.3}$$

For a definition of the thickness around the nodes $\bar{\Delta}$ see (3.6). The shape of the initial saturation can be derived from a hydrostatic profile or from the estimate at the final time of a previous data assimilation interval.

Since the initial storage is a bounded variable, the probability distribution of its error cannot conceivably be Gaussian. We therefore introduce a transform which maps the bounded variable onto the entire real axis. This transform is also an elegant way of telling the estimator that the storage is bounded, which greatly improves the estimator's convergence behavior. We map the initial storage W_g^{store} with the transform

$$\gamma_g(W_g^{\text{store}}) = \tan\left[\pi\left(\frac{W_g^{\text{store}} - W_{g,\min}^{\text{store}}}{W_{g,\max}^{\text{store}} - W_{g,\min}^{\text{store}}} - 0.5\right)\right] \tag{4.4}$$

onto a transformed initial storage γ_g varying over the entire real axis. The parameters $W_{g,\min}^{\text{store}}$ and $W_{g,\max}^{\text{store}}$ can be derived from (4.3). In practice, however, it is advisable to treat $W_{g,\min}^{\text{store}}$ and $W_{g,\max}^{\text{store}}$ as empirical parameters and to constrain the initial condition parameter W_g^{store} well within its physically possible limits. Starting a simulation at very low or very high initial saturations can entail serious numerical problems in the forward model.

4.3.2 Initial Soil Temperature and Interception Storage

The initial condition for the scaled soil temperature T_g/T_0 can be used directly. There is no need for a nonlinear transform.

Similar to the saturation, the initial interception storage must be transformed onto the entire real axis. We use the transform

$$\gamma_c(W_c) = \tan\left[\frac{\pi}{\gamma_s'}\left(\frac{W_c(t=0)}{W_c^{\max}+\gamma_s} - 0.5\right)\right] \tag{4.5}$$

Both γ_s and γ_s' are constant scaling factors.

4.3.3 Initial Condition Parameterization

The complete initial condition parameterization reads

$$\beta = \begin{bmatrix} \gamma_g^{(1)} \\ T_g^{(1)}/T_0 \\ \gamma_c^{(1)} \\ \cdots\cdots\cdots \\ \gamma_g^{(2)} \\ T_g^{(2)}/T_0 \\ \gamma_c^{(2)} \\ \cdots\cdots\cdots \\ \vdots \end{bmatrix} \qquad Y_0(\beta) = \begin{bmatrix} \begin{pmatrix} W_{g1}^{(1)}(\beta_1) \\ \vdots \\ W_{gN_z}^{(1)}(\beta_1) \end{pmatrix} \\ T_g^{(1)}(\beta_2)/T_0 \\ W_c^{(1)}(\beta_3)/W_c^{\max,(1)} + \gamma_s) \\ \cdots\cdots\cdots\cdots\cdots\cdots\cdots \\ \vdots \end{bmatrix} \tag{4.6}$$

The superscript in parentheses again indicates the number of the estimation pixel.

4.4 State Equation

Like the state vector, the state equation is a collection of equations from each pixel. The exact notation follows easily from the definition of the state vector (4.1).

The state equation for X is given by a collection of mass and energy balance equations from each pixel. In particular, the energy balance for the canopy (3.13) yields T_c, the water mass balance equation for the canopy air space (3.19) yields e_a, and the energy balance for the canopy air space (3.18) yields T_a. For a one-pixel problem, we define

$$\begin{aligned} \phi_1 &= (R_{cs}^{\text{net}} + R_{cl}^{\text{net}} - LE_{ct} - LE_{ce} - H_c)/f_c \\ \phi_2 &= LE_a - LE_g - LE_{ct} - LE_{ce} \\ \phi_3 &= H_a - H_c - H_g \end{aligned} \tag{4.7}$$

The notation for multi-pixel problems follows immediately from (4.1). Note that the canopy energy balance in ϕ_1 is divided by the fractional vegetation cover. This is important to get proper scaling of ϕ for varying vegetation densities. The mass and energy balance equations are implicit and are solved for the states with a Newton-Raphson method [Press et al., 1992].

The state equation for Y is the collection of the discretized version of Richards' equation (3.1), the force-restore equation (3.7), and the interception equation (3.15) for each

estimation pixel. For a one-pixel problem, we define

$$\begin{aligned}
\varphi_i &= \left(\sum_{j=1}^{N_z} A_{ij}\psi_j + b_j \right) / \bar{\Delta}_i / \theta_{si} \qquad \text{for} \quad 1 \le i \le N_z \\
\varphi_{N_z+1} &= \Gamma_g \left[\Gamma'_g G_g / C_g - (T_g - T_d) \right] / T_0 \\
\varphi_{N_z+2} &= (P_i - E_{ce}/\rho_w - D_c) / (W_c^{\max} + \gamma_s)
\end{aligned} \tag{4.8}$$

Again, the notation for multi-pixel problems follows immediately from (4.1). The matrix A and the vector b in the definition of φ_i, $(1 \le i \le N_z)$ follow from the discretization of Richards' equation according to the scheme by Celia et al. [1990] (see also [Simunek et al., 1997]).

$$A = \begin{pmatrix}
d_1 & e_1 & 0 & \dots & 0 \\
e_1 & d_2 & e_2 & \ddots & \vdots \\
0 & \ddots & \ddots & \ddots & 0 \\
\vdots & \ddots & e_{N_z-2} & d_{N_z-1} & e_{N_z-1} \\
0 & \dots & 0 & e_{N_z-1} & d_{N_z}
\end{pmatrix} \tag{4.9}$$

$$d = -\begin{bmatrix} 0 \\ e_1 \\ \vdots \\ e_{N_z-1} \end{bmatrix} - \begin{bmatrix} e_1 \\ \vdots \\ e_{N_z-1} \\ 0 \end{bmatrix} \tag{4.10}$$

$$e_i = (K_{ui} + K_{u,i+1})/(2\Delta_i) \qquad \text{for} \quad 1 \le i \le N_z - 1 \tag{4.11}$$

$$\begin{aligned}
b_1 &= -\bar{\Delta}_1 S_{g1} + (K_{u1} + K_{u2})/2 + q_b \\
b_i &= -\bar{\Delta}_i S_{gi} + (K_{u,i+1} - K_{u,i-1})/2 \qquad \text{for} \quad 2 \le i \le N_z - 1 \\
b_{N_z} &= -\bar{\Delta}_{N_z} S_{gN_z} - (K_{u,N_z-1} + K_{uN_z})/2 - q_t
\end{aligned} \tag{4.12}$$

4.5 Model Error

The model error or process noise accounts for errors in the model formulation, such as simplistic parameterizations or unresolved processes. In addition, it also accounts for errors in the micro-meteorologic forcings or incorrect parameter values.

In the general formulation of the estimation problem of Chapter 2, the model error is written as an additive term. Likewise, the land-surface fluxes enter the hydrologic model additively. It is therefore convenient to formulate the model error as uncertainties in the *physical* fluxes such as the latent and sensible heat fluxes. This has the added benefit that we can get a reasonable estimate of the magnitude of the model error. However, the formulation of the model error as uncertainties in the land-surface fluxes does not necessarily mean that the fluxes are the only source of uncertainty.

Formulating the model error as uncertainties in the physical fluxes determines the scaling factors D_ν and D_ω in the state equation (2.1). For a single pixel, we have $D_\nu = I_{N_X}$, where I_{N_X} is the $N_X \times N_X$ identity matrix, and

$$D_\omega(Y) = \text{diag}\left\{-1/\bar{\Delta}_1/\theta_{s1} \ldots -1/\bar{\Delta}_{N_z}/\theta_{sN_z} \quad \Gamma_g\Gamma_g'/C_g(Y)/T_0 \quad 1/(W_c^{\max}+\gamma_s)\right\} \quad (4.13)$$

where diag$\{\cdot\}$ stands for a diagonal matrix with the argument of diag$\{\cdot\}$ on the diagonal and zeros elsewhere.

The projection matrices P_ν and P_ω in the state equation (2.1) are made up of ones and zeros. They determine which fluxes are deemed uncertain. We usually restrict the model error to affect only the fluxes at the land-surface boundary. For bare soil, this means that only the flux into and out of the top node for soil moisture is uncertain. If there is vegetation, the uncertain top flux boundary condition for the saturation more directly affects all nodes that are connected to the surface by plants' roots. Owing to the rather shallow depth of the domain, all nodes are typically connected to the surface through the vegetation.

4.6 Uncertain Parameters

Land-surface hydrologic models are heavily parameterized, and few, if any, of the parameters are known accurately. The soil hydraulic parameters, for example, or the many parameters of the Radiative Transfer model are prime candidates for parameter estimation. However, estimating such parameters is no simple task. The parameter estimation problem is likely to be very ill-posed, and the benefit of estimating soil hydraulic parameters in an operational context is doubtful anyway.

At this point, we do not implement the parameter estimation in α. Our assumption is that the parameters of the hydrologic model are already well calibrated when the model is used in the data assimilation algorithm. Furthermore, we account for model deficiencies with the model error terms. But note that we do treat the parameterized initial condition as uncertain (Section 4.3).

4.7 Assimilated Data, Measurement Equation, and Downscaling

In this study, the only data assimilated into the hydrologic model are brightness temperatures. We define the data vector as

$$Z = \begin{bmatrix} T_{B1}^{\text{obs}}(t^1)/T_0 \\ \vdots \\ T_{BN_{op}}^{\text{obs}}(t^1)/T_0 \\ \cdots\cdots\cdots\cdots \\ \vdots \\ \cdots\cdots\cdots\cdots \\ T_{B1}^{\text{obs}}(t^{N_t^{\text{obs}}})/T_0 \\ \vdots \\ T_{BN_{op}}^{\text{obs}}(t^{N_t^{\text{obs}}})/T_0 \end{bmatrix} \quad (4.14)$$

The first N_{op} entries correspond to the image taken at the observation time t^1, the next N_{op} entries correspond to the image taken at the observation time t^2, etc. With a total of N_t^{obs} images and N_{op} observation pixels within each image the number of measurements is $N_Z = N_{op} \cdot N_t^{\text{obs}}$. Note that the observed brightness temperatures are scaled with the reference temperature T_0. Within the general framework of Chapter 2, it is straightforward to assimilate other types of data, for example remotely sensed surface temperatures.

Soil moisture data assimilation can be formulated as a general downscaling problem. Satellite observations of L-band brightness temperatures will likely be available at a scale much larger than the scale of the micro-meteorologic, soil texture, and land cover data. Over the continental United States, the latter data are typically available on scales of $1km$ or less, whereas the projected resolution for a space-born passive L-band sensor is on the order of $50km$ in the near future and possibly $10km$ in ten years.

Clearly, it is desirable to get soil moisture estimates at a scale finer than the resolution of the brightness images. In our formulation, the fine-scale micro-meteorologic, soil texture, and land cover data constitute valuable information incorporated into the hydrologic model. We can use such fine-scale information to effectively downscale the brightness images. By defining the measurement operator appropriately, this downscaling is easily formulated within the assimilation algorithm. In essence, we run the hydrologic model on the finer scale and define the measurement operator such that it maps the fine-scale model predictions of brightness to the coarse scale of the remotely sensed brightness data.

As an example, Figure 4.1 shows the hydrologic model defined on a scale of $5km$ (estimation pixels). By using the Radiative Transfer model (3.70) described in Section 3.2, we can derive model predictions of the brightness temperature on the $5km$-scale. In contrast, brightness data are assumed to be available on a scale of $10km$ (observation pixel). In order to compute the data misfit term $Z - M[X, Y]$, the data assimilation algorithm needs model predictions of brightness temperature on the $10km$-scale of the measurements. This is easily achieved by averaging the fine-scale brightness predictions over the observation pixels. For L-band observations, the arithmetic average is appropriate [Drusch et al., 1999; Liou et al., 1998]. In mathematical terms, we this can be summarized as

$$M_k[X, Y] \equiv \frac{1}{N} \sum_{i_k} T_{Bi_k}(t_k)/T_0 \tag{4.15}$$

The $T_{Bi_k}(t_k)$ are the predicted brightness temperatures at time t_k of the N estimation pixels i_k within the observation pixel corresponding to measurement k. We use the Radiative Transfer model (3.70) on the scale of the estimation pixels to obtain the T_{Bi_k}. Like the observations, the predicted brightness temperatures are scaled with the reference temperature T_0.

In other words, with the measurement Z_k we only supply to the algorithm the average brightness temperature over the estimation pixels within the corresponding observation pixel. But we also supply micro-meteorologic, soil texture, land cover, and other model inputs on the finer scale of the estimation pixels, that is we tell the algorithm how the estimation pixels behave relative to each other. Respecting such information through the model physics, the algorithm is then able to distribute the coarse-scale brightness update onto the finer scale estimation pixels accordingly. Note again that the downscaling procedure is carried out implicitly within the algorithm. In summary, the downscaling scheme follows naturally from an appropriate definition of the (forward) measurement operator.

4.8 Error Covariances

In principle, the formulation of the variational data assimilation algorithm allows the use of any covariance for the model and parameter errors. In practice, however, one is limited by the computational effort required to store the covariance matrices and to solve the update equations (2.16) and (2.17). Depending on the choice of covariance, the calculation of the convolution integrals for the update is easily the most computationally demanding part of the assimilation algorithm.

The computational burden of solving the update equations is greatly reduced by assuming stationarity for the error covariances. In this case, we can speed up the numerical evaluation of the convolution integral significantly by using the Fast Fourier Transform [Press et al., 1992]. Another technique to efficiently solve the convolution integral relies on solving a corresponding differential equation [Bennett et al., 1997].

It is important to note that stationarity for the model error does *not* imply that the fluxes themselves obey stationarity. For example, the moisture flux into the soil is the sum of precipitation and evaporation. Precipitation in particular is certainly not stationary in time or in space. But the stationarity assumption is only made for the *error* in the fluxes and all other contributions to the model error. The fluxes themselves are not assumed to be stationary.

We choose exponential correlation functions to model all covariances of the uncertain inputs. The initial condition parameter covariance is

$$C_\xi^{(i)(j)} = \sigma_\xi^2 \exp\left(-\frac{\sqrt{(x^{(i)}-x^{(j)})^2+(y^{(i)}-y^{(j)})^2}}{l_\xi^{xy}}\right) \tag{4.16}$$

where ξ represents the transformed initial soil moisture storage γ_g (4.4), the (scaled) initial soil temperature, or the transformed initial interception storage γ_c (4.5). The spatial coordinates of estimation pixel (i) are denoted with $x^{(i)}$ and $y^{(i)}$, and l_ξ^{xy} is the isotropic horizontal correlation length. All cross-correlations between physically different components of the uncertain parameters are assumed to vanish. For example, the error in the initial storage is always uncorrelated with the error in the initial soil temperature, regardless of the location.

For the model error, we assume that the covariance is separable in time and space. Again, physically different components of the model error are assumed to be uncorrelated. The same is assumed for saturation components belonging to different layers. For one component, the model error covariance can be written as

$$C_\xi^{(i)(j)}(t,t') = \sigma_\xi^2 \exp\left(-\frac{|t-t'|}{l_\xi^t}\right) \exp\left(-\frac{\sqrt{(x^{(i)}-x^{(j)})^2+(y^{(i)}-y^{(j)})^2}}{l_\xi^{xy}}\right) \tag{4.17}$$

where ξ represents for example the component of ν corresponding to the model error in the canopy energy balance, or the component of ω corresponding to the error in the moisture flux into and out of the top node for the saturation.

For ω, we also implemented a non-stationary model error covariance which is effectively white in time.

$$C_\xi^{(i)(j)}(t,t') = \sigma_\xi^2(t)\delta(t-t') \exp\left(-\frac{\sqrt{(x^{(i)}-x^{(j)})^2+(y^{(i)}-y^{(j)})^2}}{l_\xi^{xy}}\right) \tag{4.18}$$

The whiteness in time turns the convolution integral in the update equations into a simple multiplication. The time-dependent variance $\sigma_\xi^2(t)$ is meant to be zero most of the time, and rather big at select times, for example at times of heavy precipitation. In this way we can account for much larger errors during precipitation events without introducing too much overall uncertainty. The approach is physically motivated by the fact that quantitative precipitation measurements over large horizontal scales are usually quite uncertain (Section 7.2). Alternatively, one can choose the assimilation window in such a way that the initial condition is re-estimated after each major precipitation event (Section 7.1).

4.9 Tangent-linear and Adjoint Equations

Both the adjoint (backward) equation (2.15) and the tangent-linear (forward) model (2.18) use derivatives of the forward operators with respect to the states. It is straightforward, albeit laborious, to derive these adjoint operators from the definition of the states (4.1), the initial condition parameterization (4.6), the forward operators (4.7) and (4.8), the scaling factor $D_\omega(Y)$ (4.13), and the measurement operator (4.15). We consumed a lot of paper and pencils and accomplished this task manually. Alternatively, automatic adjoint compilers are available [Giering and Kaminski, 1998]. Since such tools have been developed only recently, they are not yet easy to use.

4.10 Temporal Discretization

So far, time has been treated as a continuum. We must, however, also discretize the time coordinate in order to implement the algorithm on a digital computer.

In the estimation algorithm, the nonlinearity of the forward model is resolved by iterating on the tangent-linear model (Section 2.2). This iteration is initialized with the prior model trajectory. Therefore we need a good initial prior model trajectory to keep the computational effort for the tangent-linear iteration at bay. For the initial prior fields, we implemented a nonlinear solver which uses a Picard iteration to handle the nonlinearity and a variable time step to deal with the intermittent nature of the forcings. The variable time step scheme follows closely the one implemented in HYDRUS, a sophisticated model for flow and transport in the unsaturated zone [Simunek et al., 1997]. The length of the time step is governed by the number of iterations needed in the Picard iteration for the most recent matric head. Moreover, the time step is reduced to seconds if there is precipitation.

In contrast, the tangent-linear model and the backward equation are implemented with a fixed time step. This time step is determined by the frequency of the micro-meteorologic inputs. Since the equations are linear and all the coefficient matrices are constant, we can solve the system directly using a time-implicit scheme. We choose the time-implicit scheme for its stability at longer time steps.

Part II

Assessing the Performance of the Assimilation Algorithm